



(12)

EUROPEAN PATENT APPLICATION

(21) Application number : **94308470.7**

(22) Date of filing : **16.11.94**

(51) Int. Cl.⁶ : **H04S 1/00, H04S 3/02, H04B 1/66**

(30) Priority : **23.11.93 US 156495**

(43) Date of publication of application :
31.05.95 Bulletin 95/22

(84) Designated Contracting States :
DE FR GB IT NL

(71) Applicant : **AT & T Corp.**
32 Avenue of the Americas
New York, NY 10013-2412 (US)


(72) Inventor : **Johnston, James David**
8 Valley View Road
Warren, New Jersey 07059 (US)
Inventor : **Sinha, Deepen**
25 Hickory Place,
Apt. B2
Chatham, New Jersey 07928 (US)

(74) Representative : **Watts, Christopher Malcolm**
Kelway, Dr. et al
AT&T (UK) Ltd.
5, Mornington Road
Woodford Green Essex, IG8 0TU (GB)

(54) **Perceptual coding of audio signals.**

(57) A five-channel perceptual audio compression system which encodes five matrixed channels switches among a fourteen encoding modes each utilizing a respective different set of matrixed channels. Six of the modes are for encoding the three front channels and eight of the modes are for encoding the two back channels. A number of the modes have as at least one of their respective matrixed channels a matrixed channel which is a function of a) one of the input channels or a sum or difference of two of them, and b) a predicted value of a).

FIG. 2



Background of the Invention

The present invention relates to the processing of information signals and, more particularly, to techniques for efficiently encoding audio signals, including signals representative of voice and music.

A significant amount of effort has been directed in recent years to so-called perceptual audio coding, or PAC. In accordance with this technique, each of a succession of time domain blocks of an audio signal is coded in the frequency domain. Specifically, the frequency domain representation of each block is divided into coder bands, each of which is individually coded, based on psycho-acoustic criteria, in such a way that the audio signal is significantly "compressed," meaning that the number of bits required to represent the audio signal is significantly less than would be the case if the audio signal were represented in a more simplistic digital format, such as in the form of PCM words.

When the audio signal comprises two or more input channels, such as the left and right channels of stereophonic (stereo) music, the above-described perceptual coding is carried out on a like number of so-called matrixed channels. In the most straightforward implementation, each matrixed channel is directly derived from a respective input channel. Thus in the stereo music case, for example, this would mean that the perceptual coding codes the frequency domain representation of the left stereo input channel over time, denoted herein as "L", and, separately, the frequency domain representation of the right stereo input channel over time, denoted herein as "R". However, further compression can be achieved when the input channels are highly correlated with one another--as, indeed, is almost always the case with stereo music channels--by switching the coding carried out for each coder band between two coding modes in which different sets of matrixed channels are used. In one of the modes, the set of two matrixed channels simply comprises the input channels L and R. In the other mode, the set of two matrixed channels comprises $S=(L+R)/2$ and $D=(L-R)/2$. The S and D channels are referred to as sum/difference channels. This technique is taught in U.S. Patent 5,285,498 issued February 8, 1994 to J. D. Johnston, hereinafter referred to as "the Johnston patent."

More recently, the art has turned its attention to the perceptual coding of more-than-two-channel audio, such as five-channel audio. (As will be apparent to those skilled in the art as this description continues, the invention can, however, be implemented in a system having other than five channels.) The input channels of a five-channel audio system typically comprise three "front" channels and two "back" channels. The front channels include the conventional left and right stereo channels plus a center channel whose frequency domain representation over time is denoted herein as C. These channels are intended to be reproduced at speakers positioned in front of the listener--at the left, at the right and directly in front, respectively. The back channels are referred to as the "left surround" and "right surround" channels whose frequency domain representations over time are denoted herein as LS and RS. These channels are intended to be reproduced at speakers positioned behind the listener--at the left and at the right, respectively.

Summary of the Invention

The above-mentioned teachings of the Johnston patent relating to switching between coding modes for the coding of stereo, i.e., two-channel audio, can be applied to a five-channel system, as well, in order, again, to provide further compression over that provided by the perceptual coding itself. For example, one can switch the front channels between two modes and the back channels between two modes. The two coding modes for the front channels would be a) a mode whose set of matrixed channels comprises L, R, and C, and b) a mode whose set of matrixed channels comprises S, D and C. Similarly, the two coding modes for the back channels would be a) a mode whose set of matrixed channels comprises LS and RS, and b) a mode whose set of matrixed channels comprises back sum/difference channels SS and SD, given by $SS=(LS+RS)/2$ and $SD=(LS-RS)/2$.

We have invented, however, a more sophisticated mode-switching approach for more-than-two-channel, e.g., five-channel, coding. In accordance with the invention, yet additional compression is achieved by switching among a plurality of modes at least one of which includes in its matrixed channel set at least one matrixed channel given by an input channel or sum/difference channel from which has been subtracted a prediction of itself. If a prediction is a "good" prediction, i.e., closely matches to the channel being predicted, the number of bits needed to represent their difference will be substantially less than that required to represent the predicted channel directly, thereby providing the aforementioned additional compression.

An example of such a mode, for the front channels, comprises the set of three matrixed channels

$$L - \hat{\alpha}C, R - \hat{\beta}C$$

where $\hat{\alpha}C$ and $\hat{\beta}C$ are predicted values of input channels L and R, respectively, as described in further detail hereinbelow. Another example of such a mode, for the back channels, comprises the set of two matrixed channels

$$SS - \hat{\eta}_1 \hat{L} - \hat{\kappa}_1 \hat{R}, \quad SD - \hat{\eta}_2 \hat{L} + \hat{\kappa}_2 \hat{R}$$

where $\hat{\eta}_1 \hat{L} - \hat{\kappa}_1 \hat{R}$ and $\hat{\eta}_2 \hat{L} + \hat{\kappa}_2 \hat{R}$ are predicted values for sum/difference channels SS and SD, respectively.

In preferred embodiments, as can be seen from the examples above, the prediction for a particular input channel or sum/difference channel is derived from at least one other, "predicting" channel. For the front channels, there are illustratively a total of six coding modes (explicitly laid out in the Detailed Description below), in which C is illustratively used as a predicting channel for L, R, S and D, while L and R are used jointly as predicting channels for C. For the back channels, there are illustratively a total of eight coding modes (also explicitly laid out below), in which all three front channels are used as predicting channels for LS, RS, SS and DS, either by themselves or in various combinations. Predictions are preferably based on the coded values of the predicting channels rather than their actual values, e.g., \hat{C} rather than C. This allows the input channels to be decoded without introducing a quantization artifact.

In preferred embodiments, the selection of which coding mode is to be used for each coder band is made by determining which of the modes will require the fewest bits to encode it.

Brief Description of the Drawing

FIG. 1 is a block diagram of a system in which the present invention is illustratively implemented;

FIG. 2 is a block diagram of the perceptual audio coder used in the system of FIG. 1;

FIG. 3 is a flowchart of a process performed within the perceptual audio coder implementing the principles of the invention;

FIGS. 4-6 illustrate a novel aspect of the perceptual audio coder relating to the use of a global masking threshold;

FIG. 7 shows the format of the PAC-encoded audio; and

FIG. 8 is a block diagram of the perceptual audio decoder used in the system of FIG. 1.

Detailed Description

To simplify the present disclosure, the following patents and publications are hereby incorporated by reference in the present disclosure as if fully set forth herein: U.S. Patent 5,040,217, issued August 13, 1991 by K. Brandenburg et al; U.S. Patent 5,341,457 issued Aug. 23, 1994 to J. L. Hall et al.; J. D. Johnston, *Transform Coding of Audio Signals Using Perceptual Noise Criteria*, IEEE Journal on Selected Areas in Communications, Vol. 6, No. 2 (Feb. 1988); International Patent Application (PCT) WO 88/01811, filed March 10, 1988; *Aspec: Adaptive Spectral Entropy Coding of High Quality Music Signals*, AES 90th Convention (1991); Johnston, J., *Estimation of Perceptual Entropy Using Noise Masking Criteria*, ICASSP, (1988); J. D. Johnston, *Perceptual Transform Coding of Wideband Stereo Signals*, ICASSP (1989); E.F. Schroeder and J.J. Platte, "MSC: Stereo Audio Coding with CD- Quality and 256 kBIT/SEC," IEEE Trans. on Consumer Electronics, Vol. CE-33, No. 4, November 1987; and Johnston, *Transform Coding of Audio Signals Using Noise Criteria*, Vol. 6, No. 2, IEEE J.S.C.A. (Feb. 1988).

For clarity of explanation, the illustrative embodiment of the present invention is presented as comprising individual functional blocks (including functional blocks labeled as "processors"). The functions these blocks represent may be provided through the use of either shared or dedicated hardware, including, but not limited to, hardware capable of executing software. Illustrative embodiments may comprise digital signal processor (DSP) hardware, and software performing the operations discussed below. Very large scale integration (VLSI) hardware embodiments of the present invention, as well as hybrid DSP/VLSI embodiments, may also be provided.

FIG. 1 is an overall block diagram of a system in which the present invention is implemented. In FIG. 1, an analog audio signal on lead 101 is fed into a preprocessor 102 where it is sampled (typically at 48 KHz) and converted into a 16-bit-per-sample digital pulse code modulation ("PCM") signal on lead 103 in standard fashion. The PCM signal is fed into a perceptual audio coder ("PAC") 104 which compresses the PCM signal and outputs the compressed PAC signal on lead 105 to either a communications channel or a storage medium 106. The latter may be, for example, a magnetic tape, compact disc or other storage medium. From the communications channel or the storage medium the compressed PAC-encoded signal on lead 107 is fed into a perceptual audio decoder 108 which decompresses the compressed PAC-encoded signal and outputs a PCM signal on lead 109 which is a digital representation of the original analog signal. From the perceptual audio decoder, the PCM signal on lead 108 is fed into a post-processor 110 which creates an analog representation.

An illustrative embodiment of the perceptual audio coder 104 is shown in block diagram form in FIG. 2. The perceptual audio coder of FIG. 2 may advantageously be viewed as comprising an analysis filter bank 202, a perceptual model processor 204, a composite coder 205, a quantizer/rate-loop processor 206 and an

entropy encoder 208.

The structure and operation of the various components of perceptual audio coder 104 are generally similar to the structure and operation of like components in FIG. 2 of the Johnston patent when processing stereo signals, and thus will not be described in detail herein except to the extent necessary for an explication of the present invention. This will include a description of composite coder 205, for which no explicit counterpart is shown in FIG. 2 of the Johnston patent (although its functionality--relating to the switching between coding modes--is implicitly carried out in the Johnston patent within analysis filter bank 202).

Turning, then, to FIG. 2 hereof, the analog audio input signal on lead 103 is illustratively a five-channel signal comprising, in the time domain, PCM samples of a set of input channels, those being left, right and center front channels and left surround and right surround back channels, denoted $l(t)$, $r(t)$, $c(t)$, $ls(t)$ and $rs(t)$, respectively. Analysis filter bank 202 receives those samples and divides them into time domain blocks. More specifically, filter bank 202 switches between two window lengths for the blocks--a "short" window comprising 128 time samples and a "long" window comprising 1024 time samples. For each block, filter bank 202 performs a Modified Discrete Cosine Transform on each of the five channels separately to provide a frequency domain representation of each channel for the block in question. The frequency domain representation of each channel includes 1024 uniformly spaced frequency spectrum lines divided into 49 coder bands for the long-window-length blocks and 128 uniformly spaced frequency spectrum lines divided into 14 coder bands for the short-window-length blocks. The frequency domain representations of the left, right, center, left surround and right surround input channels are denoted in the FIG. as $L(f)$, $R(f)$, $C(f)$, $LS(f)$ and $RS(f)$, but, for convenience, will be hereinafter denoted as L , R , C , LS and RS .

The outputs of analysis filter bank 202 are applied to composite coder 205 which in a manner described in detail below produces five matrixed channels $M_1(f)$ - $M_5(f)$ which are applied to quantizer/rate loop processor 206. The latter encodes the matrixed channels by a) generating binary data representing the signed magnitude of each of the frequency spectrum lines of each matrixed channel for a given block and b) formatting that data along with other information needed by the decoder. That data includes various pieces of "housekeeping" data as described in the Johnston patent and also hereinbelow, as well as data generated in the course of implementing the present invention including, for example, an indication of what coding mode was used to encode each coder band in a given block, as described below, as well as the values of certain prediction coefficients also described below.

The output of quantizer/rate loop processor 206 is applied to entropy encoder 208. The latter operates in conjunction with the former to achieve yet further compression.

As noted above, the output of composite coder 205 comprises five matrixed channels. Composite coder 205 has a number of coding modes each characterized by a different set of matrixed channels, the different coding modes being invoked individually for each coder band in a manner to be described. Some of the matrixed channels in some of the modes are the input channels L , R , C , LS and RS . Others of the matrixed channels are the so-called sum/difference channels $S=(L+R)/2$, $D=(L-R)/2$, $SS=(LS+RS)/2$ and $SD=(LS-RS)/2$.

In accordance with the present invention, however, at least one of the modes includes at least one matrixed channel given by an input channel or sum/difference channel from which has been subtracted a prediction of itself. In the present illustrative embodiment, six modes are provided for coding the front channels. Three, involving the front input channels L , R and C directly, are

$$\left\{ \begin{array}{ccc} L, & R, & C \\ L - \hat{\alpha}\hat{C}, & R - \hat{\beta}\hat{C}, & C \\ L, & R, & C - \hat{\gamma}\hat{L} - \hat{\delta}\hat{R} \end{array} \right\} \quad (1)$$

The other three, involving the front sum/difference channels S and D , are

$$\left\{ \begin{array}{ccc} S, & D, & C \\ S - \hat{\alpha}\hat{C}, & D - \hat{\beta}\hat{C}, & C \\ S, & D, & C - \hat{\gamma}\hat{L} - \hat{\delta}\hat{R} \end{array} \right\} \quad (2)$$

Eight modes are provided for the back channels. Four, involving the back channels LS and RS directly are

$$\left\{ \begin{array}{ll} \text{LS,} & \text{RS} \\ \text{LS} - \hat{v}_1 \hat{C}, & \text{RS} - \hat{v}_2 \hat{C} \\ \text{LS} - \hat{\eta}_1 \hat{L} - \hat{\kappa}_1 \hat{R}, & \text{RS} - \hat{\eta}_2 \hat{L} - \hat{\kappa}_2 \hat{R} \\ \text{LS} - \hat{v}_1 \hat{C} - \hat{\eta}_1 \hat{L} - \hat{\kappa}_1 \hat{R}, & \text{RS} - \hat{v}_2 \hat{C} - \hat{\eta}_2 \hat{L} - \hat{\kappa}_2 \hat{R} \end{array} \right\} \quad (3)$$

The other four, involving the back sum/difference channels SS and DS are

$$\left\{ \begin{array}{ll} \text{SS,} & \text{DS} \\ \text{SS} - \hat{v}_1 \hat{C}, & \text{DS} - \hat{v}_2 \hat{C} \\ \text{SS} - \hat{\eta}_1 \hat{L} - \hat{\kappa}_1 \hat{R}, & \text{DS} - \hat{\eta}_1 \hat{L} - \hat{\kappa}_1 \hat{R} \\ \text{SS} - \hat{v}_1 \hat{C} - \hat{\eta}_1 \hat{L} - \hat{\kappa}_1 \hat{R}, & \text{DS} - \hat{v}_2 \hat{C} - \hat{\eta}_2 \hat{L} - \hat{\kappa}_2 \hat{R} \end{array} \right\} \quad (4)$$

In the above coding modes, the terms that are subtracted from the input and sum/difference channels L, R, C, S, D, LS, RS SS and DS are, indeed, the predicted values thereof. More specifically, as can be seen, the prediction for a particular input channel or sum/difference channel is derived from at least one other, "predicting" channel. For the front channels, C is illustratively used as a predicting channel for L, R, S and D, while L and R are used jointly as predicting channels for C. For the back channels, all three front channels are illustratively used as predicting channels for LS, RS, SS and SD, either by themselves or in various combinations. The " $\hat{\cdot}$ " over the predicting channels, e.g., \hat{C} , denotes that the predictions are based on the encoded values of the predicting channels--generated by quantizer/rate loop processor 206 and fed back to composite coder 205 over lead 216--rather than their actual values. The reasoning behind this is as follows: Firstly, the encoded predicting channel serves as virtually as good a predictor as the unencoded channel so there is no harm done by doing this. It is noted however, that the predicted value of each encoded channel has to be added back in the decoder. That is, L is recovered from the transmitted $L - \hat{\alpha}\hat{C}$ by adding $\hat{\alpha}\hat{C}$ to that which is received. It is only \hat{C} , and not C, that is available in the decoder. We are thus able to add back to the received coded channel in the decoder that which was subtracted from it in the coder. If the unencoded predicting channel were used in the coder, a quantization noise artifact would, disadvantageously, be introduced in the decoder.

The scalar prediction coefficients $\hat{\alpha}$, $\hat{\beta}$, etc may be computed using a variety of different criteria. In the simplest implementation, these may all be set to "1" or some other constant less than "1". Indeed, some of those coefficients could be set to zero if the prediction of which they are a part is expected to be very small. This would be the case, for example, for the prediction of a difference matrixed channel, such as D or DS, because those matrixed channels are, themselves, expected to have near-zero values.

In a slightly more complex system, the prediction coefficients may be computed for each block (but only one set of coefficients being computed for all coder bands), using a minimum mean squared error (mmse) or a perceptually weighted mmse criterion, and transmitted as part of the bitstream that is output by quantizer/rate loop processor 206 (the composite coding decision still being independent for each of the coder bands). The reason that the coefficients are denoted with a " $\hat{\cdot}$ ", e.g., $\hat{\alpha}$, is that if their values are computed, rather than being constants, those values are necessarily coded, i.e., quantized, when stored or transmitted in digital form, the " $\hat{\cdot}$ " notation being indicative of this. In a yet more complex system, the prediction coefficients are estimated separately for each of the coder bands. Such a scheme is attractive when prediction gain is low because of time delays between different channels. However, transmitting prediction coefficients for each of the coder bands can be quite expensive. Fortunately, experiments suggest that coefficients from the previous block may be used with a relatively small loss in prediction gain. The prediction coefficients in such a system can therefore be computed in a backward fashion from the decoded values of the previous block both at the encoder and the decoder.

As seen from (1), (2), (3) and (4) above, sum/difference channels are formed only between pairs of front channels or pairs of back channels. More complex combinations, or "basis transformations",--involving perhaps three or more input channels and/or involving, perhaps, sums and differences between front and back channels--might prove advantageous from a compression point of view. Precluding such more complex basis transformations, however, is preferred because we then ensure that quantization noise masking--which is a

driving principle of perceptual coding--is effective not only in a five-speaker listening room environment but also in the case of a so-called stereo downmix of the five input channels into two channels for headphone reproduction, for example. Moreover, the more complex basis transformations will usually require the use of lower noise thresholds (discussed below) for encoding, thereby reducing the amount of compression that is achieved.

We now address the questions of a) how the encoding of the five matrixed channels of any particular coding mode is carried out, and b) how the decision is made as to which mode is to be used in order to encode a particular coder band for a given block.

In particular, the Johnston patent describes how, for each block, the perceptual model processor generates for each coder band a noise threshold for each matrixed channel, that threshold being a critical parameter for the quantizer/rate loop processor in its encoding of the respective matrixed channel. For each coder band for a two-channel system, then, four noise thresholds are available, one each for L, R, S and D. The noise thresholds are supplied by perceptual model processor 204 to composite coder 205 via lead 214. When L and R are the encoded channels for a particular coder band during one coding mode of the Johnston patent, their respective thresholds are used to encode those channels. Similarly, when S and D are the encoded channels during the other coding mode, their respective thresholds are used to encode those channels.

In the present illustrative embodiment, nine thresholds per coder band are available. They are the thresholds corresponding to the five input channels L, R, C, LS and RS and the four sum/difference channels S, D, SS and DS. Obviously, when the selected coding mode includes an input channel or sum/difference channel as one of the matrixed channels, the corresponding threshold is used. Additionally, the threshold associated with a particular input or sum/difference channel is to be used to encode each matrixed channel which includes a prediction of that input or sum/difference channel. For example, the threshold for L is also used for the matrixed channel $L - \hat{\alpha}C$. The manner in which the nine thresholds are generated for each coder band during each block is a straightforward application of the techniques described in the Johnston patent for the generation of its four thresholds and thus further elucidation on this point is not needed.

The manner in which a particular coding mode is selected is illustrated by the flowchart of FIG. 3. The flowchart represents the processing carried out in this regard for the front channels and for a particular coder band. Similar processing is carried out individually for both the front and back channels for each of the coder bands. Specifically, the items in [] relate to the processing carried out for the back channels.

Initially, the aforementioned nine thresholds are generated by perceptual model processor 204 (block 301). A decision is then made in composite coder 205 as to whether the coding mode for the front channels should be of a type that involves input channels or sum/difference channels, i.e., one of modes (1) or one of modes (2). One criterion that can be employed in making this decision is to compare the thresholds for L and R (block 303). If they differ by more than a predetermined amount, e.g., 2 dB, input channel coding is used i.e., one of the three modes in (1) (block 311). If they do not differ by more than the predetermined amount, one approach, which is not implemented here is to immediately choose sum/difference coding i.e., one of the three modes in (2). Here, however, a more sophisticated approach is used. In particular, we recognize that the use of sum/difference matrixed channels is desirable when L and R are highly correlated not only a) because a high degree of compression will be achieved in that case, but also b) because it will control so-called noise localization. Noise localization control can also be achieved, however, by suitably lowering the thresholds for L and R and, as it turns out, encoding L and R with those lowered thresholds sometimes requires fewer bits than encoding S and D. Thus one can use, for example, a "perceptual entropy" criterion as taught in the prior art to determine which approach--coding L and R with lowered thresholds or encoding S and D--will require the fewer number of bits. Once the coding mode has been narrowed to being from (1) or (2) (block 307), the particular coding mode to be used is selected simply by, again, using the aforementioned perceptual entropy criterion to identify the mode that requires the smallest number of bits (block 309 or block 314). A similar process is carried out vis-a-vis the back channels and an indication of which coding mode was used for both the front and back channels is stored or transmitted along with the encoded channels themselves.

A further novel feature of coder 104 of FIG. 2 relates to the generation of the thresholds for the five input channels and the four sum/difference channels. That feature of the coder, which uses a so-called global masking threshold, takes advantage of the masking ability of the signal component in the matrixed channel whose signal component is the strongest to mask the noise in the other matrixed channels.

In order to understand the use of this feature, we can begin by considering that in coder 104, as in the coders known in the prior art, a so-called "bit reservoir" is maintained which is basically a count of excess channel capacity measured in bits. That count, in essence, is a measure of the difference between a) the number of bit transmission slots available in the past at the average output bit rate and b) the number of bits that were actually encoded. The maximum size of the count in the bit reservoir depends on the amount of buffering (latency) allowed in the system. This unused capacity can be used to deal with the fact that future blocks may require a higher-than-average number of bits to represent the content of those blocks. Thus a relatively con-

stant output bit rate can be supported even though the number of encoded bits generated for each block varies from block to block and will, in general, be sometimes higher than that rate (on a per/second averaged basis) and sometimes lower. Typically, the bit reservoir capacity is five times the average bit rate per block.

5 If the bit reservoir reaches its maximum capacity, this is an indication that the bit requirement has been consistently lower than the average output bit rate-so much so that it is advantageous to use up such excess capacity by lowering the noise thresholds, thereby performing a finer quantization. This is advantageous in that it provides better quality reproduction. It is, however, optional, the alternative being to simply send mark, or other non-information-bearing bits.

10 A technique is, however, definitely required to deal with the case when the bit reservoir becomes empty because at that point, unless remedial steps are taken, the buffering capacity of the system will have been all used up and encoded bits will begin to be lost. The prior art deals with this problem by switching to a coarser quantization once a) the bit reservoir is depleted and, b) in addition, the bit requirement for the current block exceeds the average output bit rate per block at that time. Implementationally, this is effected by iteratively multiplying each of the noise thresholds by a series of increasing constant values greater than 1 and recalculating the bit requirement until such time as the bit requirement for the current block can be accommodated. 15 Once the bit requirements of future blocks goes below the average rate, the normal threshold values can again be used and the bit reservoir will begin to replenish.

While this prior art approach is effective in dealing with the bit reservoir depletion issue, it can lead to severe artifacts in the decoded signal since the mechanism for increasing the thresholds is deterministic across the coder bands rather than being, as we have realized is advantageous, based on psycho-acoustic considerations. 20

In accordance with the aforementioned novel feature of the coder, an approach to the bit reservoir depletion which takes psycho-acoustic considerations into account is used. That approach establishes a "global masking threshold" for each coder band which is constant across all five matrixed channels. The value of the global masking threshold is computed as described below. For the present it suffices to indicate that it represents the maximum noise level in any of the channels that will not be noticed by a listener in a listening room environment. 25

The manner in which the global masking threshold is used to control the noise threshold values in the event of imminent bit reservoir depletion is illustrated in FIGS. 4-6. FIG. 4, in particular, shows the threshold level for, illustratively, the first four coder bands of L, as well as the global masking threshold established for each of the bands. (A similar representation can be made for each of the other channels.) It is assumed that at the point in time represented by FIG. 4, the bit reservoir is just below a level of 80% depletion. Thus the standard threshold values are used. As soon as the depletion level reaches 80%, however, a lower bound for each of the thresholds of, illustratively, 50% of the global masking threshold is employed so that any threshold which is less than that lower bound is increased so as to be equal to it. Since a different global masking threshold is established for each coder band and the noise threshold is different for each band, this means that some of the thresholds will be increased more than others and some may not be increased at all. This is illustrated in FIG. 5 where it can be seen that the threshold for coder bands 1 and 3 have not been increased while those for coder bands 2 and 4 have. This approach is advantageous in that although additional noise is introduced as a consequence of increasing the thresholds, such additional noise is introduced in coder bands in which the possibility of it being detected by the listener is the lowest. The operative mechanism here is explained at a more opportune point hereinbelow. 30 35 40

If at a subsequent time it is observed that the bit reservoir reaches an even more severe state of depletion, a higher percentage of the global masking threshold, e.g. 75%, is established as the lower bound, causing a further increase in various ones of the thresholds. If, ultimately, the lower bound needs to be established at the full global masking threshold and bit reservoir depletion still continues, each of the thresholds is further increased to a value given by a constant > 1 times the relevant global masking threshold value, thereby continuing to implement psycho-acoustic considerations in the setting of the thresholds--at least to some extent. This is illustrated in FIG. 6. Once the bit reservoir begins to rebuild, lower and lower percentages of the global masking threshold come into effect for determining the noise threshold lower bounds until, at the 80% point, the thresholds return to their standard values. 45 50

The actual value of the global masking threshold for each coder band is taken to be the maximum of the five input channel thresholds for that band minus a safety margin. The safety margin, in turn, is taken to be the frequency dependent binaural masking level difference, or MLD, as defined in the Johnston patent, plus a constant 4-5 dB. 55

Given that the value of the global masking threshold is determined in this way, we are now in a position to understand why the approach described above is effective. What is happening here is that the technique takes advantage of the masking ability of the signal component in the matrixed channel whose signal compo-

nent is the strongest to mask the noise in the other matrixed channels.

There are at least two reasons why only a percentage of the global masking threshold is initially used to establish the noise threshold lower bounds, rather than the full value (unless ultimately needed as described above). One reason is that if a full global threshold is used, then the noise may not be fully masked for all listeners in a listening room environment--particularly listeners who are close to the speakers. The other is that the probability that noise in the aforementioned stereo downmixed channels will not be fully masked increases with increasing percentage of the global masking threshold used to establish the aforementioned lower bounds.

FIG. 7 shows a portion of a storage medium 700--illustratively magnetic tape--on which PAC-encoded data generated in accordance with the invention is stored and from which it can be subsequently read, decoded and presented to a listener as discussed above in connection with FIG. 1. The data is stored in frames ..., F_{i-1} , F_i , F_{i+1} , ..., each corresponding to one block of the input signal. The frames are stored sequentially on the storage medium and conform to a predetermined format that is very similar to that taught in the prior art for two-channel PAC. By way of example, the entirety of frame F_i is explicitly shown in the FIG. It has the following fields: sync word 701 defining the start of the frame; channels flag 702 indicating the number of channels in the input signal--in this example, five; window type flag 704 indicating whether the frame represents a long or short window; coding mode flags 706 which indicate which of the fourteen encoding modes was used to encode the block in question in each of the coder bands, flags 706 being represented in Huffman coded form using a predefined codebook; dc values 709 specifying a respective dc value for each of the five matrixed channels; predictor coefficient flag 710 whose value, if "1", indicates that the prediction coefficients are explicitly specified within the frame, in which case they are provided in prediction coefficient field 713 and, if "0", indicates either that all the prediction coefficients have the value 1.0 or that they are to be computed in the decoder in the manner described hereinabove, the choice between these two options being predetermined when the system is designed; and PAC data fields 714-718 each containing the encoded data for a respective one of the (in this case) five matrixed channels $M_1(f)$ through $M_5(f)$. As in prior art, two-channel PAC systems, the PAC-encoded data in each one of fields 714-718 is, in turn, represented in Huffman-coded form. Various Huffman code codebooks are used to encode the data for the various coder bands within each channel and therefore the data in each of fields 714-718 also includes information identifying which codebooks were used to encode which coder bands of the matrixed channel in question.

It will be appreciated that the format shown in FIG. 7 could equally as well be used for storing the PAC frames in other types of storage media, such as compact disc, optical disc, semiconductor memory, etc.

FIG. 8 is an illustrative embodiment of PAC decoder 109 of FIG. 1. The incoming bit stream, formatted as a sequence of frames as shown in FIG. 7, is parsed into its various components by bit stream parser 801. Although not shown explicitly in the FIG., parser 801 not only presents the PAC-encoded data to the next functional block in the decoder--entropy decoder 804-- but also supplies the various flags and other values described above in connection with FIG. 7 to each of the various components of the decoder that need them. (In similar fashion, it will be appreciated that although not explicitly shown in FIG. 2, the flags and other values that may have been generated by various components of encoder 104, to the extent that they are not generated within quantizer/rate loop processor 106 itself, are supplied thereto so as to be able to be incorporated into the format of each frame.)

Continuing with the processing carried out by decoder 109, entropy decoder 804 performs the inverse function of entropy coder 208 and provides as its outputs the five matrixed channels $M_1(f)$... $M_5(f)$. Dematrixer 807 recovers from the matrixed channels the frequency domain input channels L, R, C, LS and RS, which are denoted, as in FIG. 2, as $L(f)$, $R(f)$, $C(f)$, $LS(f)$ and $RS(f)$. The latter are then processed by Inverse Modified Discrete Cosine Transform (Inverse MDCT) processor 808 to produce the five time domain channels $l(t)$, $r(t)$, $c(t)$, $ls(t)$ and $rs(t)$.

The processing performed by each of the elements shown in FIG. 8 as just described can be implemented straightforwardly and thus needs not be described in further detail. Indeed that processing is very similar to that carried out in the prior art for two-channel PAC.

Claims

1. A method of perceptually encoding an audio signal comprising a set of input channels, the method comprising the steps of
 - a) generating a set of matrixed channels in response to said input channels, and perceptually encoding said matrixed channels, characterized in that at least an individual one of the matrixed channels of said set of matrixed channels is a function of a) an individual one of said input channels or the sum of, or the difference between,

two of them, and b) a prediction of a).

2. The invention of claim 1 comprising the further step of alternatively applying to an output a) said perceptually encoded matrixed channels and b) a perceptual encoding of said input channels.
- 5 3. The invention of claim 1 comprising the further step of alternatively applying to a communication channel or a storage medium a) said perceptually encoded matrixed channels and b) a perceptual encoding of said input channels.
- 10 4. The invention of claim 1 wherein in said audio signal, said input channels are represented in the frequency domain and wherein said generating step comprises the step of generating said set of matrixed channels from said input channels in such a way that each of said matrixed channels is represented, for each of a series of time domain blocks, by frequency spectrum lines.
- 15 5. The invention of claim 4 wherein said frequency spectrum lines are divided into a plurality of coder bands and wherein the magnitudes of the frequency spectrum lines of each coder band are represented by values that are quantized as a function of a noise threshold associated with that coder band.
- 20 6. Apparatus for processing an audio signal comprising a set of input channels, said apparatus comprising means (205) for perceptually encoding one of a predetermined plurality of sets of matrixed channels generated in response to said input channels, and means (206, 208) for applying the perceptually encoded matrixed channels to a selected one of a) a communications channel, and b) a storage medium, characterized in that the channels of an individual one of said sets of matrixed channels are said input channels and the channels of others of said sets of matrixed channels include at least an individual one matrixed channel which is a function of a) an individual one of said input channels or the sum of, or the difference between, two of them, and b) a prediction of a).
- 25 7. The invention of claim 6 wherein in said audio signal, said input channels are represented for each of a series of time domain blocks, by frequency spectrum lines divided into a plurality of coder bands, the magnitudes of the frequency spectrum lines of each coder band being represented by values that are quantized as a function of a noise threshold associated with that coder band.
- 30 8. A method for processing a perceptually encoded audio signal, said perceptually encoded audio signal having been generated by generating a set of matrixed channels in response to a set of input channels; perceptually encoding said matrixed channels; and applying said perceptually encoded matrixed channels to a communication channel or a storage medium, said method comprising the steps of receiving said perceptually encoded matrixed channels from said communications channel or storage medium, decoding the received perceptually encoded matrixed channels, and recovering said input channels from the decoded matrixed channels, characterized in that said set of matrixed channels comprises a selected one of i) said input channels, and ii) a set of matrixed channels in which at least an individual one of the matrixed channels is a function of a) an individual one of said input channels or the sum of, or the difference between, two of them, and b) a prediction of a).
- 35 9. The invention of claim 8 wherein said recovering step includes the step of determining how said set of matrixed channels was generated in response to said set of input channels.
- 40 10. Apparatus for processing a perceptually encoded audio signal, said perceptually encoded audio signal having been generated by generating a set of matrixed channels in response to a set of input channels; perceptually encoding said matrixed channels; and applying the perceptually encoded matrixed channels to a communication channel or a storage medium; said apparatus comprising means (801) for receiving said perceptually encoded matrixed channels from said communications channel or storage medium, means (804) for decoding the received perceptually encoded matrixed channels, and means (807) for recovering said input channels from the decoded matrixed channels,
- 45 50 55

characterized in that at least an individual one of the matrixed channels of said set of matrixed channels being a function of a) an individual one of said input channels or the sum of, or the difference between, two of them, and b) a prediction of a).

- 5 11. The invention of claim 10 wherein said means for recovering determines how said set of matrixed channels was generated in response to said set of input channels.
12. An article (700) in which is stored information representing a perceptually encoded audio signal comprising a set of input channels, said perceptually encoded audio signal having been generated by the steps of
10 of
generating a set of matrixed channels in response to said input channels, and
perceptually encoding said matrixed channels,
characterized in that at least an individual one of the matrixed channels of said set of matrixed channels is a function of a) an individual one of said input channels or the sum of, or the difference between,
15 two of them, and b) a prediction of a).
13. The invention of claim 12 wherein said perceptually encoded audio signal was generated by the further step of perceptually encoding said input channels alternatively with said matrixed channels.
- 20 14. The invention of claim 1 or claim 6 or claim 8 or claim 10 or claim 12 wherein said individual one of said matrixed channels is a function of the difference between a) and b).
15. The invention of claim 1 wherein said set of input channels includes left, right and center channels.
- 25 16. The invention of claim 15 wherein a first one of said matrixed channels is a function of said left channel and a prediction of said left channel; wherein a second one of said matrixed channels is a function of said right channel and a prediction of said right channel; and wherein a third one of said matrixed channels is said center channel; each said prediction being a function of said center channel.
- 30 17. The invention of claim 16 wherein each said prediction is a function of a perceptually encoded version of said center channel.
18. The invention of claim 16 wherein each said prediction is a function of the product of a) a perceptually encoded version of said center channel, with b) a respective prediction coefficient.
- 35 19. The invention of claim 15 wherein a first one of said matrixed channels is said left channel; wherein a second one of said matrixed channels is said right channel; and wherein a third one of said matrixed channels is a function of said center channel and a prediction of said center channel, that prediction being a function of said left and right channels.
- 40 20. The invention of claim 19 wherein said prediction is a function of perceptually encoded versions of said left and right channels.
21. The invention of claim 19 wherein said prediction is a function of the product of a) a perceptually encoded version of said left channel with b) a respective prediction coefficient, and is further a function of the product of a) a perceptually encoded version of said right channel, with b) a respective prediction coefficient.
- 45 22. The invention of claim 15 wherein a first one of said matrixed channels is a function of the sum of said left and right channels; wherein a second one of said matrixed channels is a function of the difference between said left and right channels; and wherein a third one of said matrixed channels is said center channel.
- 50 23. The invention of claim 15 wherein a first one of said matrixed channels is a function of a sum channel and a prediction of said sum channel; wherein a second one of said matrixed channels is a function of a difference channel and a prediction of said difference channel; and wherein a third one of said matrixed channels is said center channel; said sum channel being a function of the sum of said left and right channels, said difference channel being a function of the difference between said left and right channels, and
55 each said prediction being a function of said center channel.
24. The invention of claim 23 wherein each said prediction is a function of a perceptually encoded version of

said center channel.

25. The invention of claim 23 wherein each said prediction is a function of the product of a) a perceptually encoded version of said center channel, with b) a respective prediction coefficient.
- 5 26. The invention of claim 15 wherein a first one of said matrixed channels is a sum channel; wherein a second one of said matrixed channels is a difference channel; and wherein a third one of said matrixed channels is a function of said center channel and a prediction of said center channel; said sum channel being a function of the sum of said left and right channels, said difference channel being a function of the difference between said left and right channels, and said prediction being a function of said left and right channels.
10
27. The invention of claim 26 wherein said prediction is a function of perceptually encoded versions of said left and right channels.
- 15 28. The invention of claim 26 wherein said prediction is a function of the product of a) a perceptually encoded version of said left channel with b) a respective prediction coefficient, and is further a function of the product of a) a perceptually encoded version of said right channel, with b) a respective prediction coefficient.
- 20 29. The invention of claim 15 wherein said set of input channels further includes left surround and right surround channels.
30. The invention of claim 29 wherein an individual one of said matrixed channels is a function of said left surround channel and a prediction of said left surround channel, and wherein a further one of said matrixed channels is a function of said right surround channel and a prediction of said right surround channel.
25
31. The invention of claim 30 wherein said left surround channel prediction and said right surround channel prediction are respective functions of one or more of said left, right and center channels.
32. The invention of claim 30 wherein said left surround channel prediction and said right surround channel prediction are respective functions of perceptually encoded versions of one or more of said left, right and center channels.
30
33. The invention of claim 29 wherein an individual one of said matrixed channels is a function of the sum of said left surround and right surround channels; and wherein a further one of said matrixed channels is a function of the difference between said left surround and right surround channels.
35
34. The invention of claim 29 wherein an individual one of said matrixed channels is a function of a sum surround channel and a prediction of said sum surround channel; and wherein a further one of said matrixed channels is a function of a difference surround channel and a prediction of said difference surround channel; said sum surround channel being a function of the sum of said left surround and right surround channels, and said difference surround channel being a function of the difference between said left surround and right surround channels.
40
35. The invention of claim 34 wherein said sum surround channel prediction and said difference surround channel prediction are respective functions of one or more of said left, right and center channels.
45
36. The invention of claim 34 wherein said sum surround channel prediction and said difference surround channel prediction are respective functions of perceptually encoded versions of one or more of said left, right and center channels.

50

55

FIG. 1




FIG. 2

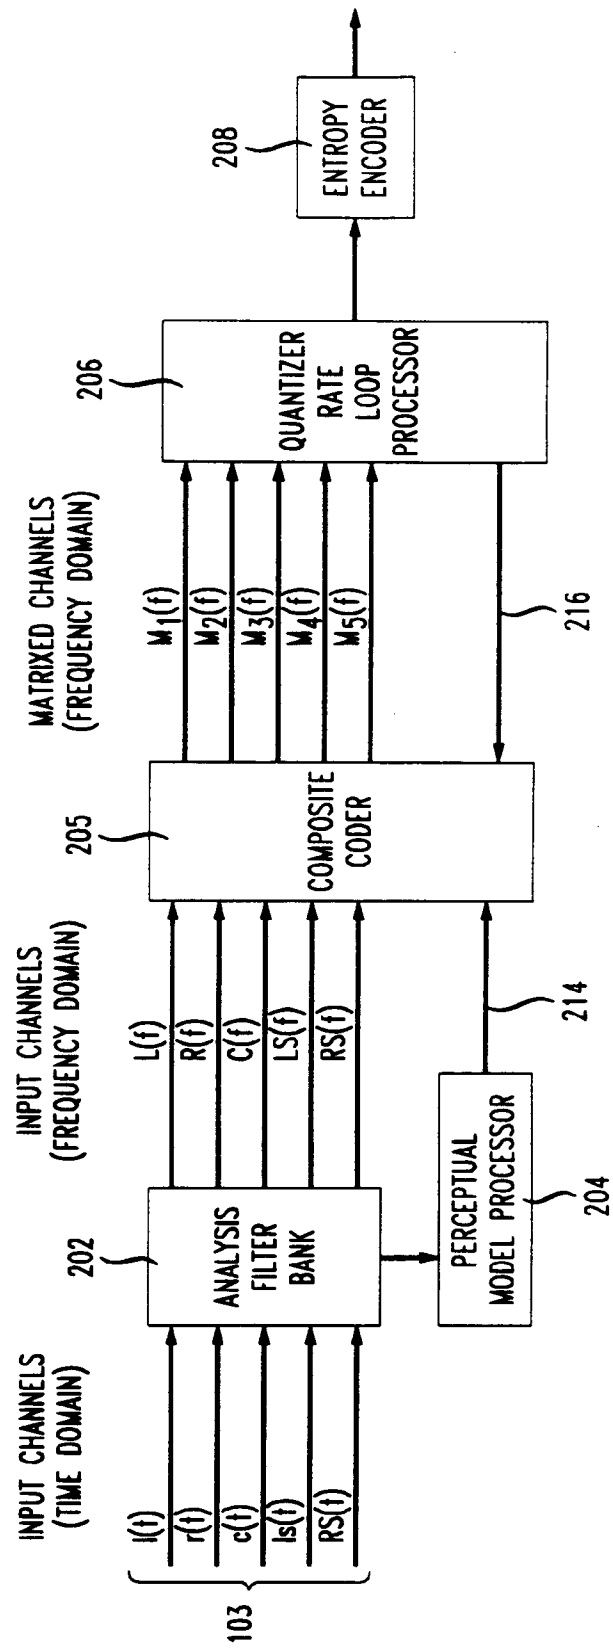


FIG. 3

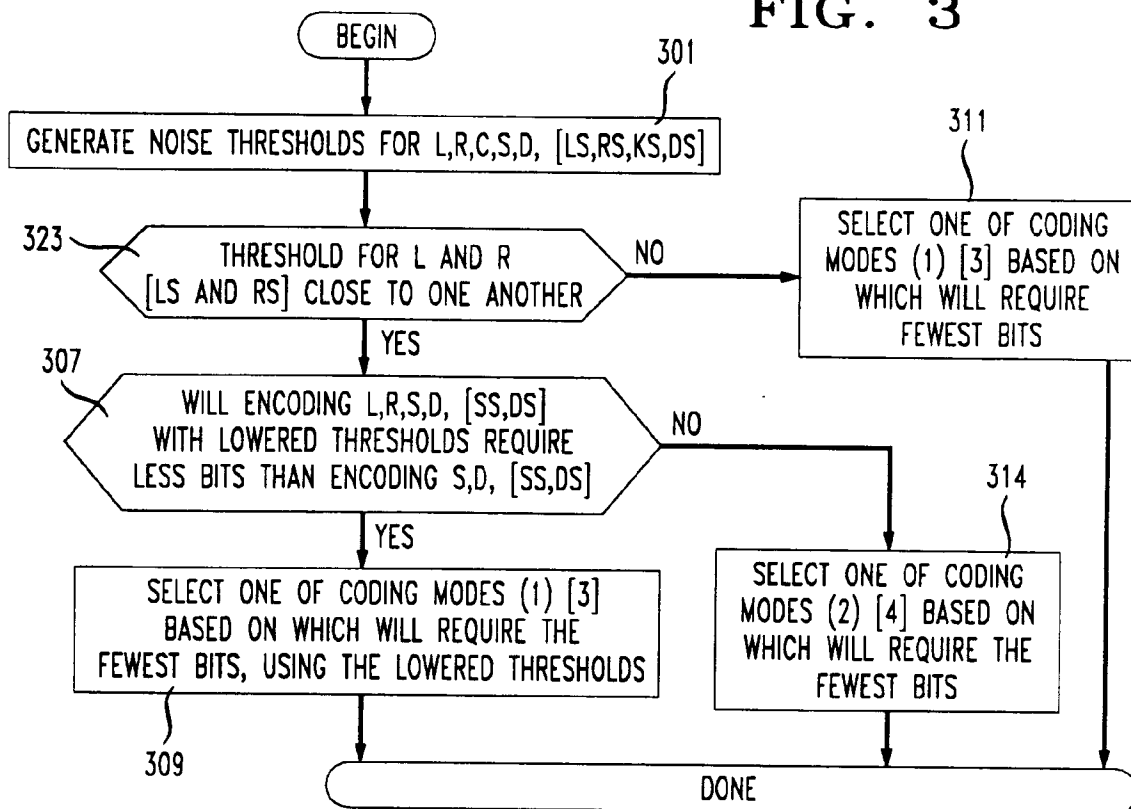


FIG. 4




FIG. 6




FIG. 7




FIG. 5





FIG. 8





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 94 30 8470

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	EP-A-0 564 089 (AT & T) * page 2, line 11-13 * * page 3, line 24-52 * * page 4, line 40 - page 11, line 38 * * page 13, line 54 - page 17, line 39 * ---	1-3,5-8, 10,12, 13,15, 17,18, 20,21, 24,25, 27-29, 32,36	H04S1/00 H04S3/02 H04B1/66
D,A	IEEE J.S.C.A., vol. 6, no.2, February 1988 page 314-323 J.D.JOHNSTON 'TRANSFORM CODING OF AUDIO SIGNALS USING PERCEPTUAL NOISE CRITERIA.' * paragraph I-IV * ---	1-3,5-8, 10,12, 13,15, 17,18, 20,21, 24,25, 27-29, 32,36	TECHNICAL FIELDS SEARCHED (Int.Cl.6) H04S H04B H04H H04N
A	IEEE ICASSP-92, vol. 2, 23 March 1992 SAN FRANCISCO, pages 205-208, XP 000356973 KATE ET AL. 'MATRIXING OF BIT RATE REDUCED AUDIO SIGNALS.' * page 205, left column, line 1 - page 208, left column, line 16 * ---	1-15,19, 22,29,33	
A	WO-A-90 16136 (BRITISH TELECOMMUNICATION P.L.C.) * page 1, line 1-3 * * page 3, line 32 - page 4, line 10 * * page 7, line 13 - page 12, line 35 * * page 13, line 8 - page 17, line 26 * -----	1,6,8, 10,12, 14-36	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 28 February 1995	Examiner Zanti, P
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>I : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons ----- & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.82 (P/MCO1)